



Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure

J. Patrick Gary
Network Projects Leader/606
NASA Goddard Space Flight Center



GODDARD SPACE FLIGHT CENTER




Science Driver

- New NASA Science Needing Gigabit per Second (Gbps) Networks
 - Coordinated Earth Observing Program
 - Hurricane Predictions
 - Global Aerosols
 - Remote viewing & Manipulation of Large Earth Science Data Sets
 - Integration of Laser and Radar Topographic Data with Land Cover Data
 - Large-Scale Geodynamics Ensemble Simulations
- Advances in Networking Technology
 - National LambdaRail (NLR) implementation
 - Global Lambda Integrated Facility (GLIF) cooperation

Next Step: OptIPuter, NLR, and Starlight Enabling Coordinated Earth Observing Program (CEOP)

Source: Milt Halem, NASA GSFC

Accessing 300TB's of Observational Data in Tokyo and 100TB's of Model Assimilation Data in MPI in Hamburg -- Analyzing Remote Data Using GRaD-DODS at These Sites Using OptIPuter Technology Over the NLR and Starlight



**Note Current Throughput 15-45 Mbps:
OptIPuter 2005 Goal is ~10 Gbps!**



<http://ensight.eos.nasa.gov/Organizations/ceop/index.shtml>



OptIPuter and NLR will Enable Daily Land Information System Assimilations

- **The Challenge:**
 - More Than Dozen Parameters at ~ 50 GB per Parameter, Produced Six Times A Day, Need to be Analyzed
- **The LambdaGrid Solution:**
 - Sending this Amount of Data to NASA Goddard from Project Columbia at NASA Ames for Human Analysis Would Require < 15 Minutes/Day Over NLR
- **The Science Result:**
 - Making Feasible Running This Land Assimilation System Remotely in Real Time



Source: Milt Halem, NASA GSFC





NLR/GSFC Applications: Hurricane Prediction

- The NASA Finite-Volume General Circulation Model (fvGCM) has been producing real-time, high-resolution (~25 km) weather forecasts focused on improving hurricane track and intensity forecasts.
- During the active 2004 Atlantic hurricane season, the fvGCM provided landfall forecasts with an accuracy of ~100 km up to 5 days in advance.
- The 50–100 Mbps throughput available between fvGCM users at GSFC and the Columbia supercomputer at ARC greatly hindered carrying out time-critical simulations of the hurricanes that devastated Florida.
- The 10 Gbps NLR access will enable remote, 3D visualization analysis as soon as forecast variables become available.
- Key Contacts: Ricky Rood, Bob Atlas, Horace Mitchell, GSFC; Chris Henze, ARC.



In an fvGCM forecast, Hurricane Frances makes landfall on the Gulf Coast of Florida while Hurricane Ivan intensifies in the tropical Atlantic. Visualization by J. Williams, GST.




NLR/GSFC Applications: Global Aerosols

- Project Atmospheric Brown Clouds (ABC) is an international effort to discover and analyze areas of brown colored atmosphere to learn how dust and pollution particles are transported and what impacts they have on the environment, climate, agricultural cycles, and quality of life.
- GSFC and the Scripps Institution of Oceanography (SIO) are planning a collaboration to predict the flow of aerosols from Asia across the Pacific to the U.S. on timescales of days to a week.
- GSFC will provide an aerosol chemical tracer model (GOCAR) embedded in a high-resolution regional model (MM5) that can assimilate data from Indo-Asian and Pacific ground stations, satellites, and aircraft.
- Remote computing and analysis tools running over the NLR will enable acquisition and assimilation of the Project ABC data.
- Key Contacts: Yoram Kaufman, William Lau, GSFC; V. Ramanathan, Chul Chung, SIO.



Strategically located ground stations in the Indo-Asian and Pacific regions monitor atmospheric pollution.




The global nature of brown clouds is apparent in analysis of NASA MODIS Data. Research by V. Ramanathan, C. Corrigan, and M. Ramana, SIO.



NLR/GSFC Applications: Remote Viewing and Manipulation of Large Earth Science Data Sets

- Remote viewing and manipulation of data sets at GSFC and JPL is needed to support EOSDIS and Earth system modeling.
- GSFC's EOSDIS Clearing House (ECHO) and JPL's GENESIS prototype science analysis system (iEarth) will become connected over the NLR. The link will enable comparison of hundreds of terabytes of data, generating large, multi-year climate records.
- Initial work will focus on the Estimating the Circulation and Climate of the Ocean (ECCO) modeling team. Besides ready access to the NLR, the team will need versatile subsetting and other data manipulation functions to reduce compute and bandwidth requirements as well as a set of Grid-accessible statistical analysis and modeling operators to refine and validate the ECCO models.
- Key Contacts: ECHO metadata gateway team, GSFC; GENESIS team, led by Tom Yunck, JPL.




Near-surface (15-m) ocean current speed from an eddy-permitting integration of the cubed-sphere ECCO ocean circulation model. Research by JPL and MIT. Visualization by C. Henze, Ames.





NLR/GSFC Applications: Integration of Laser and Radar Topographic Data with Land Cover Data

- NASA has executed two advanced missions to create an accurate high-resolution topographic model of the Earth: the Shuttle Radar Topography Mission (SRTM) and ICESat, with its Geoscience Laser Altimeter System (GLAS).
- The agency now has the opportunity to merge the two data sets, using SRTM to achieve good coverage and GLAS to generate calibrated profiles. Proper interpretation requires extracting land cover information from Landsat, MODIS, ASTER, and other data archived in multiple DAACs.
- Use of the NLR and local data mining and subsetting tools will permit systematic fusion of global data sets, which are not possible with current bandwidth.
- Key Contacts: Bernard Minster, SIO; Tom Yunck, JPL; Dave Harding, Claudia Carabajal, GSFC.




<http://icesat.gsfc.nasa.gov>

<http://www2.jpl.nasa.gov/srtm>

<http://glcf.umiacs.umd.edu/data/modis/vcf>

APPLICATIONS - Future GRID on 10-GE Network



Dr. Zhou is working on applying Grid Computing and High-Speed Network to large-scale distributed computing in Earth and Space Science. More details can be found at <http://esto.nasa.gov/conferences/estc2004/papers/a4p1.pdf>.

NLR Will Provide an Experimental Network Infrastructure for U.S. Scientists & Researchers



“National LambdaRail” Partnership
Serves Very High-End Experimental and Research Applications
4 x 10Gb Wavelengths Initially
Capable of 40 x 10Gb wavelengths at Buildout



First Light
September 2004



Task Objective

- “...establish a “Lambda Network” (in this case using optical wavelength technology and 10 Gbps Ethernet per wavelength) from GSFC’s Earth science Greenbelt facility in MD to the Scripps Institute of Oceanography (SIO) through the University of California, San Diego (UCSD) facility over the National Lambda Rail (NLR), a new national dark optical fiber infrastructure.”
- “...make data residing on Goddard’s high speed computer disks available to SIO with access speeds as if the data were on their own desktop servers or PC’s.”
- “...enable scientists at both institutions to share and use compute intensive community models, complex data base mining and multi-dimensional streaming visualization over this highly distributed, virtual working environment.”



Accomplishments for the Year

- Partner with NSF-funded OptIPuter Project - national leaders in optical WAN networking, distributed cluster computing, and mega-pixel visualization display research
 - Early 10-GE connection with NLR/CAVEwave lambda
 - Free use of 10-Gbps WASH-STAR lambda
 - OptIPuter networking with Scripps Institute of Oceanography
- Partner with NSF-funded DRAGON Project - national leaders in optical MAN networking research
 - Two 10-Gbps and three 2.4-Gbps lambdas initially, of 40 possible
- Access to Other 10-Gbps NLR lambdas: Shared IP, GE VLANs, HOPI
- First 10-Gbps network within GSFC
- Leading NASA's way in NLR use for ARC's Project Columbia



NASA GSFC Among First 10 Users of the NLR

- **GSFC's initial 10-Gbps connection to the NLR was enabled via cooperation with the National Science Foundation (NSF)-funded OptIPuter Project (<http://www.optiputer.net>)**
- **GSFC's initial 10-Gbps NLR connection was used to transmit Earth science data sets in real time to an OptIPuter 15-screen tiled display at the SC2004 conference in Pittsburgh, PA.**
- **“The involvement of NASA Goddard demonstrated the capabilities of NLR and showed just how researchers in ‘big science’ will need this kind of capacity to make new discoveries about aspects of our world and to help transfer this knowledge to practical uses by others in carrying out important tasks that improve our lives.”**
 - **Tom West, President and CEO of the NLR**

NASA GSFC in the NLR booth with the OptIPuter-provided 15-screen tiled display cluster during SC2004



NLR booth at SC2004 with OptIPuter-provided 15-screen tiled display cluster.



Eric Sokolowsky (GST, Inc.) of GSFC's SVS interactively views model and observation data (set 1) from NASA's Animated Earth project with hyperwall paradigm.



Eric Sokolowsky (GST, Inc.) of GSFC's SVS with model and observation data (set 2) from NASA's Animated Earth project in hyperwall paradigm.



Randall Jones (GST, Inc.) of GSFC's SVS with model data from NASA's Land Information System in OptIPuter's display paradigm.



Various visitors to the NLR booth being briefed by Tom West, president and CEO of the NLR.




Rear view of the OptIPuter-provided 15-screen tiled display cluster.

L-Net SC2004 Photo Gallery: <http://esdc.gsfc.nasa.gov/LNetphoto3.html>


Photo Sources: Randall Jones, NASA GSFC

NASA GSFC Tests with OptIPuter Across the National LambdaRail

January 2005



GSFC L-Net Configurations at McLean and Greenbelt





Future Work

- **MAP Core Integration LambdaGrid Infrastructure**
 - New science drivers and evaluators of NLR interconnection among USCD/SIO, UIC, GSFC, JPL, ARC
 - Coordinated Earth Observing Program
 - Hurricane Predictions
 - Global Aerosols
 - Remote viewing & Manipulation of Large Earth Science Data Sets
 - Integration of Laser and Radar Topographic Data with Land Cover Data
 - Collaboration among PI Larry Smarr (UCSD/Cal-(IT)2), Co-I's John Orcutt (UCSD/SIO), Tom DeFanti (UIC), Milt Halem (UMBC), and several scientists at GSFC, JPL, & ARC
- **High-Speed Networking, Grid Computing, and Large-Scale Ensemble Simulations in Geodynamics, Weijia Kuang (GSFC), Shujia Zhou (GSFC) et al**
- **Expanding 10-GE L-Net**
 - More science buildings/clusters within GSFC; More NLR dedicated lambdas, e.g: ARC, ORNL, GISS; Wide Area SAN for NCCS; Optical switching within GSFC



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Special Acknowledgements

GSFC Internal

- **IT Pathfinder Working Group**
 - Chair: Dr. Milton Halem/Emeritus & UMBC
 - Applications Lead: Mike Seablom/610.3
 - Middleware Lead: Walt Truszkowski/588
 - Network Lead: Pat Gary/606.1

- **High End Computer Network Team**
 - Bill Fink/606.1
 - Kevin Kranacs/585
 - Paul Lang/ADNET/606.1
 - Aruna Muppalla/ADNET/606.1
 - Jeff Martz/CSC/606.2
 - Mike Steffenelli/CSC/606.2
 - George Uhl/SWALES/423
 - Steve Booth/SWALES/423
 - Kevin Fisher/586/UMBC coop

GSFC External

- **National LambdaRail**
 - CEO: Tom West
 - Net Eng Lead: Debbie Montano

- **OptIPuter Project (NSF-funded)**
 - PI: Dr. Larry Smarr/UCSD
 - Co-PI: Dr. Tom DeFanti/UIC
 - PM: Maxine Brown/UIC
 - UCSD Net Eng: Greg Hidley, Arron Chin, Phil Papodopolos
 - UCIC Net Eng: Alan Verlo, Linda Winkler

- **DRAGON Project (NSF-funded)**
 - PI: Jerry Sobieski/UMCP
 - Co-I: Tom Lehman/USC-ISI/E
 - Net Eng: Chris Tracy

- **NASA Research and Education Network**
 - DPM: Kevin Jones/ARC



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Principal Investigator & Co-Investigators

- **Name:** Pat Gary (930) & Jeff Smith (585) Co-PI's & GSFC's Information Technology Pathfinder Working Group (ITPWG) as Co-I's
- **Organizations:** Code 420, Code 580, Code 920, & Code 930
- **Telephone:** 301-286-9539 & 301-614-5038 for Co-PI's
- **E-mail:** Pat.Gary@nasa.gov, JeffSmith@nasa.gov for Co-PI's

Project Website

- http://esdcd.gsfc.nasa.gov/IRAD_Lambda.html



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Backup Slides





GSFC FY04 IRAD Project "Preparing Goddard for Large Scale
Team Science in the 21st Century: Enabling an All Optical
Goddard Network Cyberinfrastructure"

**End of Year Review for
GSFC Technology Management Office
February 18, 2005**

J. Patrick Gary
Network Projects Leader
Computational & Information Sciences and Technology Office/606
NASA Goddard Space Flight Center
pat.gary@nasa.gov
301-286-9539





GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Outline for End of Year Review

- **Motivation**
 - **Advances in Networking Technology**
 - **Enabling New NASA Science Needs**
- **Goals**
- **Key Challenges and Solution Designs**
- **Implementation Status**
- **Next Steps**




GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Motivation Outline

- **Advances in Networking Technology**
 - Bandwidth growth rate greater than Tflops growth rate
 - National LambdaRail (NLR) implementation
 - Global Lambda Integrated Facility (GLIF) cooperation
 - Latest Internet2 IPv4 Land Speed Record
 - Personal Computer Interface
- **New NASA Science Needing Gigabit per Second (Gbps) Networks**
 - Coordinated Earth Observing Program (CEOP)
 - Hurricane Predictions
 - Global Aerosols
 - Remote viewing & Manipulation of Large Earth Science Data Sets
 - Integration of Laser and Radar Topographic Data with Land Cover Data
 - Large-Scale Geodynamics Ensemble Simulations

Optical WAN Research Bandwidth Has Grown Much Faster than Supercomputer Speed!



Source: Timothy Lance, President, NYSERNet



NLR Will Provide an Experimental Network Infrastructure for U.S. Scientists & Researchers



“National LambdaRail” Partnership
Serves Very High-End Experimental and Research Applications
4 x 10Gb Wavelengths Initially
Capable of 40 x 10Gb wavelengths at Buildout



**First Light
September 2004**

For more information regarding NLR see <http://www.nlr.net> or contact info@nlr.net

Internet2 Land Speed Record

(Rules and current records: <http://lsr.internet2.edu/>)

Latest IPv4 Single Stream Record (<http://data-reservoir.adm.s.u-tokyo.ac.jp/lsr>)

- **7.21 Gbps** (TCP payload), standard frame, 148.850 Petabit meter / second
- 20,645 km connection between SC2004 booth and CERN through Tokyo, Latency 433 ms RTT



Network used in the experiment




Personal Computer Interface (PCI) Advances

- **Shared Parallel Bus**
 - PCI 1.0 (32-bit, 33 MHz): 1.056 Gbps (1 direction at a time)
 - PCI 2.3 (64-bit, 66 MHz): 4.224 Gbps (1 direction at a time)
 - PCI-X 1.0 (64-bit, 133 MHz): 8.448 Gbps (1 direction at a time)
 - PCI-X 2.0 (64-bit, 266 MHz): 16.896 Gbps (1 direction at a time)
- **Dedicated Serial Interface (4 wires per “lane”)**
 - **PCI Express:**
 - 2.5 Gbps (raw) per lane each direction
 - 2.0 Gbps (without encoding overhead) per lane each direction (maximally 4.0 Gbps bi-directional)
 - Up to 32 lanes

Next Step: OptIPuter, NLR, and Starlight Enabling Coordinated Earth Observing Program (CEOP)

Source: Milt Halem, NASA GSFC

Accessing 300TB's of Observational Data in Tokyo and 100TB's of Model Assimilation Data in MPI in Hamburg -- Analyzing Remote Data Using GRaD-DODS at These Sites Using OptIPuter Technology Over the NLR and Starlight



**Note Current Throughput 15-45 Mbps:
OptIPuter 2005 Goal is ~10 Gbps!**



<http://ensight.eos.nasa.gov/Organizations/ceop/index.shtml>



OptIPuter and NLR will Enable Daily Land Information System Assimilations

- **The Challenge:**
 - More Than Dozen Parameters at ~ 50 GB per Parameter , Produced Six Times A Day, Need to be Analyzed
- **The LambdaGrid Solution:**
 - Sending this Amount of Data to NASA Goddard from Project Columbia at NASA Ames for Human Analysis Would Require < 15 Minutes/Day Over NLR
- **The Science Result:**
 - Making Feasible Running This Land Assimilation System Remotely in Real Time



Source: Milt Halem, NASA GSFC





NLR/GSFC Applications: Hurricane Prediction

- The NASA Finite-Volume General Circulation Model (fvGCM) has been producing real-time, high-resolution (~25 km) weather forecasts focused on improving hurricane track and intensity forecasts.
- During the active 2004 Atlantic hurricane season, the fvGCM provided landfall forecasts with an accuracy of ~100 km up to 5 days in advance.
- The 50–100 Mbps throughput available between fvGCM users at GSFC and the Columbia supercomputer at ARC greatly hindered carrying out time-critical simulations of the hurricanes that devastated Florida.
- The 10 Gbps NLR access will enable remote, 3D visualization analysis as soon as forecast variables become available.
- Key Contacts: Ricky Rood, Bob Atlas, Horace Mitchell, GSFC; Chris Henze, ARC.



In an fvGCM forecast, Hurricane Frances makes landfall on the Gulf Coast of Florida while Hurricane Ivan intensifies in the tropical Atlantic. Visualization by J. Williams, GST.




NLR/GSFC Applications: Global Aerosols

- Project Atmospheric Brown Clouds (ABC) is an international effort to discover and analyze areas of brown colored atmosphere to learn how dust and pollution particles are transported and what impacts they have on the environment, climate, agricultural cycles, and quality of life.
- GSFC and the Scripps Institution of Oceanography (SIO) are planning a collaboration to predict the flow of aerosols from Asia across the Pacific to the U.S. on timescales of days to a week.
- GSFC will provide an aerosol chemical tracer model (GOCAR) embedded in a high-resolution regional model (MM5) that can assimilate data from Indo-Asian and Pacific ground stations, satellites, and aircraft.
- Remote computing and analysis tools running over the NLR will enable acquisition and assimilation of the Project ABC data.
- Key Contacts: Yoram Kaufman, William Lau, GSFC; V. Ramanathan, Chul Chung, SIO.



Strategically located ground stations in the Indo-Asian and Pacific regions monitor atmospheric pollution.




The global nature of brown clouds is apparent in analysis of NASA MODIS Data. Research by V. Ramanathan, C. Corrigan, and M. Ramana, SIO.



NLR/GSFC Applications: Remote Viewing and Manipulation of Large Earth Science Data Sets

- Remote viewing and manipulation of data sets at GSFC and JPL is needed to support EOSDIS and Earth system modeling.
- GSFC's EOSDIS Clearing House (ECHO) and JPL's GENESIS prototype science analysis system (iEarth) will become connected over the NLR. The link will enable comparison of hundreds of terabytes of data, generating large, multi-year climate records.
- Initial work will focus on the Estimating the Circulation and Climate of the Ocean (ECCO) modeling team. Besides ready access to the NLR, the team will need versatile subsetting and other data manipulation functions to reduce compute and bandwidth requirements as well as a set of Grid-accessible statistical analysis and modeling operators to refine and validate the ECCO models.
- Key Contacts: ECHO metadata gateway team, GSFC; GENESIS team, led by Tom Yunck, JPL.




Near-surface (15-m) ocean current speed from an eddy-permitting integration of the cubed-sphere ECCO ocean circulation model. Research by JPL and MIT. Visualization by C. Henze, Ames.



NLR/GSFC Applications: Integration of Laser and Radar Topographic Data with Land Cover Data

- NASA has executed two advanced missions to create an accurate high-resolution topographic model of the Earth: the Shuttle Radar Topography Mission (SRTM) and ICESat, with its Geoscience Laser Altimeter System (GLAS).
- The agency now has the opportunity to merge the two data sets, using SRTM to achieve good coverage and GLAS to generate calibrated profiles. Proper interpretation requires extracting land cover information from Landsat, MODIS, ASTER, and other data archived in multiple DAACs.
- Use of the NLR and local data mining and subsetting tools will permit systematic fusion of global data sets, which are not possible with current bandwidth.
- Key Contacts: Bernard Minster, SIO; Tom Yunck, JPL; Dave Harding, Claudia Carabajal, GSFC.



<http://icesat.gsfc.nasa.gov>


<http://www2.jpl.nasa.gov/srtm>

<http://glcf.umiacs.umd.edu/data/modis/vcf>



GODDARD SPACE FLIGHT CENTER

APPLICATIONS - Future GRID on 10-GE Network



Dr. Zhou is working on applying Grid Computing and High-Speed Network to large-scale distributed computing in Earth and Space Science. More details can be found at <http://esto.nasa.gov/conferences/estc2004/papers/a4p1.pdf>.



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Project Goals

- "...establish a "Lambda Network" (in this case using optical wavelength technology and 10 Gbps Ethernet per wavelength) from GSFC's Earth science Greenbelt facility in MD to the Scripps Institute of Oceanography (SIO) through the University of California, San Diego (UCSD) facility over the National Lambda Rail (NLR), a new national dark optical fiber infrastructure."
- "...make data residing on Goddard's high speed computer disks available to SIO with access speeds as if the data were on their own desktop servers or PC's."
- "...enable scientists at both institutions to share and use compute intensive community models, complex data base mining and multi-dimensional streaming visualization over this highly distributed, virtual working environment."




GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Key Challenges and Solution Designs Outline (1 of 2)

- **Implementing 10-Gbps Computer Networks End-to-End (ISO Layers 1-3)**
 - **Transcontinental Network Part**
 - **NLR Phase 1/Year 1**
 - **Regional Network Part**
 - **DRAGON Phase 1/Year 1**
 - **Local Area Network Part**
 - **10-GE upgrade to GSFC's Scientific and Engineering Network**

NLR Phase 1 - Installation Schedule

Will Complete Aug 2004




NLR Wavelengths




- **Initial complement of 4 λ s installed and available at outset**
 - One λ for national switched Ethernet experimental network
 - Another λ for national 10 Gbps IP network to support internetworking and end-to-end transport protocol experiments
 - Similar to Internet2's Abilene except routers will be available for measurement and experimentation
 - Third λ will serve as a quick start facility for new research projects
 - Fourth λ will be used by Internet2's HOPI testbed
- More λ s will be activated as needed to support the research and operational objectives of the community**

Dynamic Resource Allocation with GMPLS on Optical Networks (DRAGON) Configuration



GSFC L-Net Configurations at McLean and Greenbelt





GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Key Challenges and Solution Designs Outline (2 of 2)

- **Tuning Applications for High Performance Networks Use (ISO Layers 4-7)**
 - Large round-trip-time latencies for packet acknowledgements
 - TCP Alternates or Enhancements
 - Slow disk access times
 - Pre-fetch caching to RAM
 - Interactive data steaming to 100 mega-pixel displays
 - Multiple GE interfaces to visualization clusters
 - GrADS/DODS
 - Porting to OptIPuter connected hosts




GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Outline for End of Year Review

- **Motivation**
- **Goals**
- **Key Challenges and Solution Designs**
- **Implementation Status**
- **Next Steps**



DRAGON eVLBI Experiment Configuration





NASA GSFC Among First 10 Users of the NLR

- **GSFC's initial 10-Gbps connection to the NLR was enabled via cooperation with the National Science Foundation (NSF)-funded OptIPuter Project (<http://www.optiputer.net>) and the NLR (<http://www.nlr.net>)**
- **GSFC's initial 10-Gbps NLR connection was used to transmit Earth science data sets in real time to an OptIPuter 15-screen tiled display at the SC2004 conference in Pittsburgh, PA.**
- **“The involvement of NASA Goddard demonstrated the capabilities of NLR and showed just how researchers in ‘big science’ will need this kind of capacity to make new discoveries about aspects of our world and to help transfer this knowledge to practical uses by others in carrying out important tasks that improve our lives.”**
 - **Tom West, President and CEO of the NLR**

NASA GSFC in the NLR booth with the OptIPuter-provided 15-screen tiled display cluster during SC2004

- **Earth science data sets created by GSFC's Scientific Visualization Studio were retrieved across the NLR in real time and displayed at the SC2004 in Pittsburgh**
- **Animated Earth (<http://aes.gsfc.nasa.gov/>) data sets were retrieved from OptIPuter servers in Chicago and San Diego and from GSFC servers in McLean, VA**
- **Land Information System (<http://lis.gsfc.nasa.gov/>) data sets were retrieved from OptIPuter servers in Chicago, San Diego, & Amsterdam**



NLR booth at SC2004 with OptIPuter-provided 15-screen tiled display cluster. Photo Source: Randall Jones, NASA GSFC

L-Net SC2004 Photo Gallery
<http://esdcd.gsfc.nasa.gov/LNetphoto3.html>

Interactive Retrieval and Hyperwall Display of Earth Sciences Images on a National Scale

Enables Scientists To Perform Coordinated Studies Of Multiple Remote-Sensing Or Simulation Datasets

Source: Milt Halem & Randall Jones, NASA GSFC & Maxine Brown, UIC EVL




Earth science data sets created by GSFC's Scientific Visualization Studio were retrieved across the NLR in real time from OptIPuter servers in Chicago and San Diego and from GSFC servers in McLean, VA, and displayed at the SC2004 in Pittsburgh



<http://esdcd.gsfc.nasa.gov/LNetphoto3.html>





Randall Jones

NASA GODDARD SPACE FLIGHT CENTER



Global 1 km x 1 km Assimilated Surface Observations Analysis
Remotely Viewing ~ 50 GB per Parameter



NASA GSFC in the NLR booth with the OptIPuter-provided 15-screen tiled display cluster during SC2004



NLR booth at SC2004 with OptIPuter-provided 15-screen tiled display cluster.



Eric Sokolowsky (GST, Inc.) of GSFC's SVS interactively views model and observation data (set 1) from NASA's Animated Earth project with hyperwall paradigm.



Eric Sokolowsky (GST, Inc.) of GSFC's SVS with model and observation data (set 2) from NASA's Animated Earth project in hyperwall paradigm.



Randall Jones (GST, Inc.) of GSFC's SVS with model data from NASA's Land Information System in OptIPuter's display paradigm.



Various visitors to the NLR booth being briefed by Tom West, president and CEO of the NLR.




Rear view of the OptIPuter-provided 15-screen tiled display cluster.

L-Net SC2004 Photo Gallery: <http://esdcd.gsfc.nasa.gov/LNetphoto3.html>

Photo Sources: Randall Jones, NASA GSFC

NASA GSFC Tests with OptIPuter Across the National LambdaRail

January 2005






NASA GSFC Among First 10 Users of the NLR

- **Presently GSFC's computers connected to the NLR are located in the NLR suite at the Level3 Communications' optical fiber "carrier hotel" facility in McLean, VA**
- **In early March of 2005, two 10-Gbps connections will be enabled across the NSF-funded multi-wavelength Dynamic Resource Allocation via GMPLS Optical Network (DRAGON) research network (<http://dragon.east.isi.edu>)**
- **These DRAGON-based connections will link NLR/McLean with several high-performance computers at GSFC's main site in Greenbelt, MD, as well as with computers at other sites on the Washington, DC-area DRAGON**
- **Access to other 10-Gbps NLR lambdas is planned via membership in Mid-Atlantic Terascale Partnership (for the Shared IP and GE VLAN lambdas) and participation in Internet2's Hybrid Optical and Packet Infrastructure**




DRAGON Network

Washington, DC Metro Area




MATP Aggregation Facility Architecture

DRAFT



- 10 gigE or OC192
- 1 gigE
- Expansion not limited to number of lines shown
- WDM

HOPi Node





GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Outline for End of Year Review

- **Motivation**
- **Goals**
- **Key Challenges and Solution Designs**
- **Implementation Status**
- **Next Steps**







High Performance Remote Data Access Via GSFC L-Net Follow-ons

- **Extending 10-GE L-Net within GSFC to more science buildings/clusters**
- **Dedicated 10-GE NLR lambda(s) between GSFC and:**
 - NASA ARC
 - UCSD/SIO & OptIPuter
 - ORNL
 - UIC/OptIPuter
- **GISS on shared or dedicated 10-GE NLR lambda**
- **Wide Area SAN: CXFS-SGI between NAS and NCCS**
- **Optical switch for both GSFC's East and West campuses**

High Performance Networking and Remote Data Access GSFC L-Net for NCCS and Science Buildings

Legend

- Dark Fiber
- 10 Gbps GE
- 2 Gbps FC
- 1 Gbps GE





GSFC L-Net Enabling New NASA Science Needs

- **New science drivers and evaluators of NLR interconnection among UCSD/SIO, UIC, GSFC, JPL, ARC**
 - Coordinated Earth Observing Program
 - Hurricane Predictions
 - Global Aerosols
 - Remote viewing & Manipulation of Large Earth Science Data Sets
 - Integration of Laser and Radar Topographic Data with Land Cover Data
- Reference: "MAP Core Integration LambdaGrid Infrastructure" proposal, January 14, 2005
- PI: Larry Smarr (UCSD/Cal-(IT)2)
 - Co-I's: John Orcutt (UCSD/SIO), Tom DeFanti (UIC), Milt Halem (UMBC)
- **W. Kuang et al., "High Speed Networking and Large-Scale Simulation in Geodynamics", abstract/poster, Fall AGU 2004**
 - **S. Zhou et al., "High-Speed Network and Grid Computing for High-End Computation: Application in Geodynamics Ensemble Simulations", submitted for 13th Annual Mardi Gras Conference, February 2005**



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Major Significance (1 of 2)

- **Partner with NSF-funded OptIPuter Project**
 - Collaboration with national leaders in optical WAN networking, distributed cluster computing, and mega-pixel visualization display research
 - Early 10-GE connection with NLR/CAVEwave lambda
 - Free use of 10-Gbps WASH-STAR lambda
 - OptIPuter networking with Scripps Institute of Oceanography
- **Partner with NSF-funded DRAGON Project**
 - Collaboration with national leaders in optical MAN networking research
 - Two 10-Gbps and three 2.4-Gbps lambdas initially, of 40 possible
- **Access to Other 10-Gbps NLR lambdas**
 - Shared IP and GE VLANs via membership in Mid-Atlantic Terascale Partnership
 - Internet2's Hybrid Optical and Packet Infrastructure



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Major Significance (2 of 2)

- **First 10-Gbps network within GSFC: inter- and intra-buildings connecting with science user compute/storage/visualization clusters**
- **Enabling new NASA science needs**
 - **Coordinated Earth Observing Program (CEOP)**
 - **Hurricane Predictions**
 - **Global Aerosols**
 - **Remote viewing & Manipulation of Large Earth Science Data Sets**
 - **Integration of Laser and Radar Topographic Data with Land Cover Data**
 - **Large-Scale Geodynamics Ensemble Simulations**
- **Leading the way in NLR use for ARC's Project Columbia**



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Special Acknowledgements

GSFC Internal

- **IT Pathfinder Working Group**
 - Chair: Dr. Milton Halem/Emeritus & UMBC
 - Applications Lead: Mike Seablom/610.3
 - Middleware Lead: Walt Truszkowski/588
 - Network Lead: Pat Gary/606.1
- **High End Computer Network Team**
 - Bill Fink/606.1
 - Kevin Kranacs/585
 - Paul Lang/ADNET/606.1
 - Aruna Muppalla/ADNET/606.1
 - Jeff Martz/CSC/606.2
 - Mike Steffenelli/CSC/606.2
 - George Uhl/SWALES/423
 - Steve Booth/SWALES/423
 - Kevin Fisher/586/UMBC coop

GSFC External

- **Nationa LambdaRail**
 - CEO: Tom West
 - Net Eng Lead: Debbie Montano
- **OptIPuter Project (NSF-funded)**
 - PI: Dr. Larry Smarr/UCSD
 - Co-PI: Dr. Tom DeFanti/UIC
 - PM: Maxine Brown/UIC
 - UCSD Net Eng: Greg Hidley, Arron Chin, Phil Papodopolos
 - UCIC Net Eng: Alan Verlo, Linda Winkler
- **DRAGON Project (NSF-funded)**
 - PI: Jerry Sobieski/UMCP
 - Co-I: Tom Lehman/USC-ISI/E
 - Net Eng: Chris Tracy
- **NASA Research and Education Network**
 - DPM: Kevin Jones/ARC



GSFC FY04 IRAD Project "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Principal Investigator & Co-Investigators

- **Name:** Pat Gary (930) & Jeff Smith (585) Co-PI's & GSFC's Information Technology Pathfinder Working Group (ITPWG) as Co-I's
- **Organizations:** Code 420, Code 580, Code 920, & Code 930
- **Telephone:** 301-286-9539 & 301-614-5038 for Co-PI's
- **E-mail:** Pat.Gary@nasa.gov, Jeff.Smith@nasa.gov for Co-PI's

Project Website

- http://esdcd.gsfc.nasa.gov/IRAD_Lambda.html